Feide system architecture

December 2007

English translation: Apr. 15th, 2008

Architect: Ingrid Melve, Cato Olsen

Technical writer: Ketil Albertsen

Document history: Version 1.0

Table of contents

1 Introduction

This document describes the Feide architecture.

1.1 Target audience

The primary target audience is IT management personell who needs to identify actors and components in the Feide federation, and the interaction between these at an overview level. No detail technical knowledge is required, but some familiarity with IT systems is an advantage.

Technical details are described in *Feide Login Service Requirement Specification* [1]**Error! Reference source not found.** . The present report makes references to specific elements of these requirements.

1.2 What is understood by an architecture

The Feide architecture is comprised of:

- A model of components, information and concepts of Feide.
- Specification and textual description of the *components* interacting to realize the Feide federation.
- Specification and textual description of the *interfaces* between components in the Feide federation.
- Contracts, regulations and recommendations to which actors in the Feide federation must adhere.

1.3 Symbols used in the illustrations

Illustrations use Feides standardized symbols to represent the architectural components:

👤 - Feide users: Students, pupils, parents, employees

● - Feide services: Learning systems, libraries, information vendors

⌂ - Host organizations: Educational, research and similar institutions

ⓘ - Feide and its login service Moria

∩ - Other federations: eduroam, MinSide (the Norwegian public information federation)

2 Feide survey

2.1 Feide as identity provider

Feide, through its login service Moria, provides in Feide, med innloggingstjenesten Moria, mediates information about Feide users to service providers.

Conforming that a user is the person he claims to be is called to *authenticate* that Feide user. A person using Feide services is authenticated once by *logging in* through Moria at the start of a working session. Throughout the session, Feide will attest the identity of the user to the various service providers; Feide is a *trusted third party.*

Feide also communicaties reliable information, *user attributes,* regarding authenticated Feide users. A service may therefore be relieved of the tasks of managing basic user data, and of keeping these data up to date.

Authentication confirms the user's identity, independent of the rights, the *authorization*, that user is entitled to. The service may determine the user's authorization based on user attributes communicated by Feide. The authorization may be directly given by the attribute values, e.g. the user's organizational affiliation or kind of affiliation (student, employee). The service may, if it knows the identity of the user, manage its own authorization information at an individual level.

Feide offers authentication and user information based on web protocols, and may be used by services offered across the web. The technical solutions employed in the current implementation is not adapted to systems with user interfaces based on other technologies.

2.2 Feide organizations

Feide manages exchange of information between actors managing *Feide users*, *host organizations*, and actors offering services to users, *service providers*. A given organization, such as a university, may act both as a host organization and as a service provider. As seen by Feide, these are distinct functions with differing requirements. Any organization having signed a contract with Feide is called a *Feide organization.*

Host organizations, service providers, and Feide together makes up a *federation.*

Figur 1: Feide federation overview

The user employs a web browser for communicating with a number of web based services. The services receive authentication information and user attributes through Feide. This information is not saved by the Feide login service, but are retrieved from *LDAP directories* in the host organization of the user. User attributes are accessed with the access privileges to the directory given to the user himself, and Feide mediates the attributes to the services requiring them. Before login, the user may interrogate the attributes Feide may access about himself, and approve the exposure of these to the service to be activated, in accordance with the contract between Feide and the service provider. If the user does not want these attributes to be exposed, he may cancel the login.

2.3 Feide login

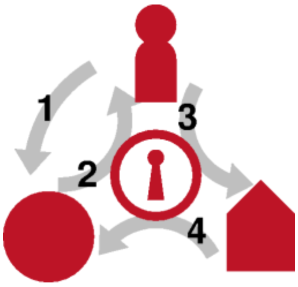

A successful Feide login goes through the following steps:

1. The user attempts to open a web page for the service he wants to activate
2. The service makes an authentication request to Feide, and Feide displays a login form to the user.
3. The user fills in his Feide user name and password in the login form and returns it to Feide. Feide forwards the user name and password to the user's host organization for verification.
4. The host organization returns to Feide a confirmation that the user is authenticated, along with the attributes stored for this user in the host organization's LDAP directory. Feide forwards to the service a confirmation of the authentication of the user, along with the selection of user attributes the service is entitled to according to the contract with Feide.

If a user has already specified his Feide name and password when activating another service, this sequence of operations may be shortcut: When Feide receives the authentication request in step 2, a confirmation of the user's identity along with relevant attributes may be returned to the service without requesting the Feide name and password from the user, and the host organization need not be contacted.

2.4 Federating Feide users and services

Making a Feide user known to a service provider is called to *federate* the user with the service. In the primary application environments of Feide, the user is first established as a Feide user through a host organization, and later federated with a service provider account. If the service has no need to recognize the Feide user from one session to another, a *one-time federation* may be automatically established at the start of the session, to be dissolved at the end of the session.

When desired or required, information about several users may be transferred in advance from a host organization to a service provider by *provisioning:* In a single operation, the host organization uploads information about e.g. all new students this semester, to the service. The upload is performed independently of the Feide services. When the user makes his first connect to the service, Feide conveys sufficient information about the user to allow the service to identify the appropriate (pre-registered) account for federation with the user. Provisioning is of particular interest when host organization and service provider are the same organization (i.e. local services), and the service requires information not defined in the Feide LDAP scheme.

2.5 Authentication through other federations, cross federating

Some users are not affiliated with any Feide organization, but are known in other federations offering single sign on and similar services. These users may be given access to Feide services by cross federating. One example is parents of elementary school pupils; the parents may be granted access to Feide services by logging in to the Norwegian public authentication service «MinSide».

Service providers not belonging to Feide may also recognize Feide authentication for users of their own services, such as Norwegian researchers gaining access to the supercomputing facilities of the Finish universities through Feide authentication.

Feide cooperates with several federations to allow a user authenticated by one federation to use the services offered by another federation without having to go through a second login procedure. Cross federating requires the cooperating federations to trust each other's authentication procedures. When a service requests authentification of a user from another federation, and attribute information for this user, Feide forwards the request to the federation responsible for authenticating the user.

2.6 Information forwarded through Feide vs. local information

A Feide user identifies himself by his *Feide name*. In the user interface, the user specifies his local user name and selects his organizational affiliation from a drop down list. Internally, these are combined to one string: <UserName>@<HostOrganization>. The second part indicates where the Feide login service should go for authentification and information, *attributes*, about the user. A host organization manages a standard set of attributes for each user, defined by a *Feide scheme*, which is a Nordic adaptation of the international standard schemes eduPerson and eduOrg [3]. The Nordic adaptation is referred to as *norEdu* [2].

The user attributes are stored in an *LDAP directory* of the host organisation the user belongs to. When a user is authenticated, Feide retrieves information from the host organization's LDAP directory and forwards the selection of attributes that the service is entitled to according to the contract with Feide. This relieves the service from managing information about individual users: Rather than maintaining its own attribute store, the service asks Feide to provide the information on demand. E.g. a library need not keep track of the email addresses of its users; when a user places a book reservation, the library is informed by Feide where to send a notification (i.e. the user's email address) when the book becomes available.

A service receiving information through Feide will always receive fully updated information. However, some services have a need for more information than what is available through Feide, and must locally maintain its own supplementary information directory. Usually, it is beneficial to use the local directory for supplementary attributes only, avoiding duplicate storing of information that is available through Feide. This will ensure that the service does not rely on possibly outdated information.

Feide fully controls which information is forwarded from host organization to service provider. Contracts with each individual service provider limits the user attributes made available: If a service provider cannot demonstrate a need for knowing the identity of the person requesting the service, no information about that person is revealed. Feide may still indicate e.g. that the user is an (unidentified) student. Information is revealed only when required to perform the service, and in agreement with the host organizations manging this information.

Feide stores no user attributes beyond the temporary storing necessary while a Feide session is active. The login service maintains a persistent directory of federation keys for each service a user is federated with.

2.7 Identification of individuals

Basically, the user's Feide name is unknown to a service provider, which only knows a *federation key* which provides no information about the user. The contract with Feide may specify that a selection of attributes are to be transferred at authentication time, when required to perform the service. The federation key may be used to identify e.g. an individua profile archived by the service provider, a local account, or for local access control. The user's federation key is different for each service, so information from different services regarding a given user cannot be correlated based on the federation key.

Some attributes, e.g. the National Identification Number, does provide unique identification of a person, and may be used to correlate information from different sources. Certain other attributes, such as email addresses, provide a near-unique identification of a person. Feide follows a restrictive policy in exposing attributes that may identify individuals, and service providers must demonstrate a real need in order to access these through Feide.

2.8 Using the previous generation of Feide, Moria 2

The present document describes the Moria III software which was put into operation in 2007. It implies minimal changes to the end user , but because the new architecture adopts standard protocols, the systems are not compatible at a technical level. All new services will employ the new protocols.

Feide will support Moria 2 authentification through summer of 2008. Service providers are encouraged to switch to the new solution well before Moria 2 is taken out of service. Technical information about Moria 2 is available at moria.sourceforge.net [10].

3 The Feide information model

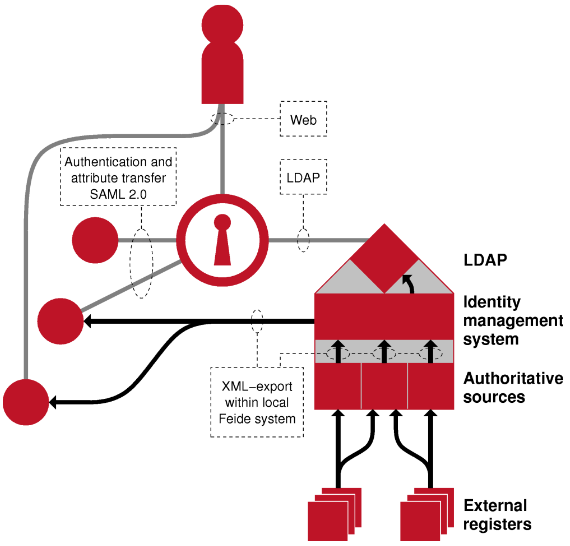
Web
Authentication and attribute transfer SAML 2.0
LDAP
LDAP
Identity management system
Authoritative sources
XML–export within local Feide system
External registers

Figure 2: Information sources and Identity Management System

The information flow in Feide runs from host organizations, through Feide, to service providers. Any information flow from service providers to host organizations bypasses Feide and is not defined in this architecture.

Feide makes certain requirements to the information management in host organization taking part in the federation, and the host organization is committed to honor the instructions specified in the contract with Feide. To ensure that information is structured, correct and up to date at any time, Feide requires each host organization to implement an automated information management system. An important component in this structure is an automatic *user management system*, UMS, as shown in the figure above.

A UMS handles information, retrieved from authoritative sources, regarding the organization's users. Examples are contact information, affiliation type (student, employee,...), authentification information etc.

Feide requires each host organization to provide a standard set of attribute values describing each of the organization's users, in an LDAP directory containing copies of the relevant parts of the information handled by the UMS of the organization. Semantics and structure of the information is specified in the Feide LDAP scheme [2], which is based on international cooperation and current de-facto-standarder: the eduPerson/eduOrg-schemes developed by EDUCAUSE [3].

The LDAP directory has the following properties:

- The LDAP directory is employed for user authentification.
- The LDAP directory contains descriptive attributes for all authenticated users.
- The information in the LDAP directory is, as far as possible, kept up to date and correct at all times. Detail requirements is indicated in guidelines specified in the contract between Feide and the host organization.
- The login service retrieves user information from the LDAP directory of the user's host organization.

The following requirements must be satisfied for all LDAP directories:

- The communication between the directory and Feide is reliable and protected against eavesdropping.
- The directory is available through the standard LDAP protocol.
- The directory information is supplied from a user management system (a UMS) of the organization.
- The login service must be informed about the DN («Distinguished Name») of the root of the Feide part of the directory tree, and must be given sufficient access rights to search for a user's DN based on his Feide name.

For a formal specification, see *Feide Login Service Requirement Specification* [1], chapter 4, requirements 26 to 36.

4 The Feide network of trust

The Feide network of trust is expressed through a set of agreements, contracts, implementations and guidelines. The term «network of trust» is an indication that a certain level of mutual confidence between the actors is required for the federation to work across organizational borders.

The current implementation of Feide defines a single level of trust.

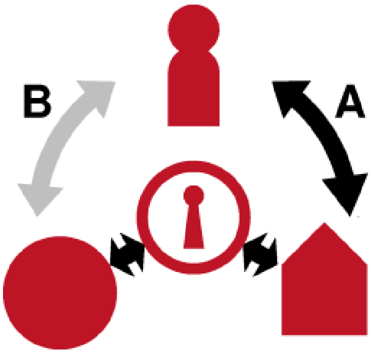

Figure 3: Implicit and explicit trust relations

Feide users are managed by their host organizations. Feide requires that the users are held responsible for all their actions when accessing computer systems and services, and that agreed rules for acceptable use are enforced.

Requirement for protection of information and privacy is regulated by Norwegian law. When Feide authentication is used with host organizations or service providers outside of Norway, similar laws in other countries may apply. Feide has been developed to satisfy the requirements in the Norwegian Privacy Protection Act [6], with particular attention to security and protection of information about individuals.

The arrows in figure 3 above indicates trust relationships. Employees, students and others belonging to an organization have a trust relationship to their host organization, shown by the arrow «A». Modifications of this relationship over time is managed by the host organization. Because Feide's contact with users goes through host organizations, requirements and obligations to Feide will affect users indirectly only.

The trust relationship between users and services, labelled «B» in Figure 3, is an implicit relationship that users may choose to accept by logging in to Feide. The trust is based on the security mechanisms of Feide, the contract between the service provider and the host organization, Feide's agreements with these two, and between the user an his host organization.

Several aspects of the mutual responsibilities and rights between the federation actors are managed through technical devices. Yet, several aspects rely on appropriate implementation of procedures and guidelines, where one party must trust the other party to act according to these.

Below is identified those conditions that are specified in the form of contracts, as well as areas where cooperation is based on a confidence that the contract is lived up to.

4.1 The contract between Feide and a host organization

The contract between Feide and a host organization sets the following conditions:

- technical details regarding the host organization's LDAP directory,
- legal aspects of responsibility between Feide and the host organization,
- procedures to ascertain confidentiality and privacy protection,
- availability, format and quality of user information submitted to Feide,
- responsibility for user support.

Obligations of Feide:

- provide 2nd line support,
- continuously monitoring the operation of the login service,
- make arrangements for monitoring access to the LDAP directory.

Obligations of the host organization

- provide authentication of its own users through the LDAP directory of the organization,
- provide user attributes according to the syntax and semantics defined by the Feide LDAP scheme,
- deliver consistent and up to date information to the Feide login service,
- provide 1st line support to its own users.

The host organization trusts Feide to:

- correctly identify a service provider accessing user attributes through Feide,
- protect personal information, including password, against unauthorized access, and not to use this information for other purposes than authorizing Feide users and forwarding of user attributes to service providers according to their contracts.

Feide trusts the host organization to:

- correctly confirm or reject an assertion of the password of one of the organization's users,
- on request deliver correct user attributes according to contract,
- reject authentification of users who are no longer affiliated with the organization.

4.2 Contract between Feide and a service provider

The contract between Feide and a service provider sets the following conditions:

- technical details regarding the service,
- legal responsibilities of Feide and of the service provider,
- a list of host organization to be given access to the service; the service may also be open to all host organizations or only for users affiliated with the service provider's organization,
- which attributes will be forwarded to the service when a user has been authenticated. Feide offers a selection of predefined alternatives: an anonymous profile, a small set of limited profiles, and a complete profile.
- contacts for management and technical issues.

Obligations of Feide

- enable Feide login to local services for users affiliated with the service provider's organization,
- inform all host organizations when a new common service is introduced,
- ensure that users from host organizations that have accepted the common service, and no other users, are granted access to the service.

Obligations of the service provider:

- file all Feide enabled services, both local services and common services, with UNINETT,
- maintain a personal information protection statement which must be easily accessible by Feide users.

The service provider trusts Feide to:

- deliver correct information about the identity of authenticated users.

Feide trusts the service provider to:

- desist from permanently storing user attributes delivered by Feide, except when as specified in contract when required for providing the service.

4.3 Cross federation

The Feide federation may cross federate with other federations, through agreements that sets the following conditions:

- The direction of federation: Whether externally authenticated users shall be allowed access to Feide services, and/or Feide authenticated user shall be allowed access to services in the other federation.
- Legally responsible persons representing each federation.
- Which technical standards the cross federation is based on.
- Contacts for administrative and technical issues.
-
- Guidelines for attribute exchange: Which attributes may be exchanged (this depends on the direction of federation) and how attributes in other federations are mapped to Feide attributes.

Feide will only cross federate with other federations that authenticate users at a confidence level comparable to Feide's own procedures.

4.4 Handling of external risks

Feide risikoanalyse [8] («Feide risk analysis», in Norwegian) surveys a number of threats against confidentiality, reliability and access, and identifies technical protection measures where relevant, trust relationships where relevant.

5 Interfaces

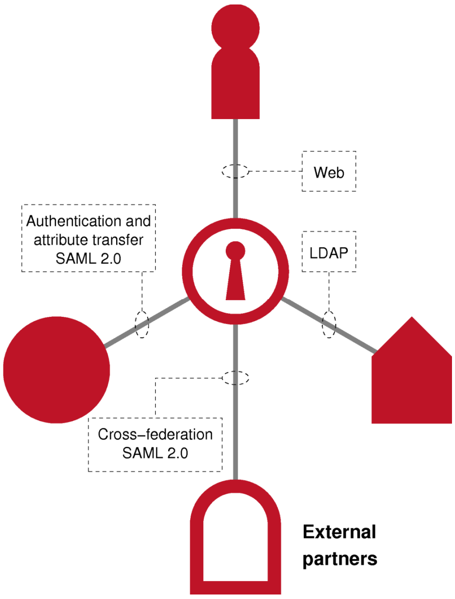

Figure 4: Components and interfaces in the Feide federation

The figure above illustrates the following interfaces:

- *Web:* Between end users and the Feide login service.
- *Authentication and attribute forwarding:* Between Feide services and the Feide login service.
- *Cross federation:* Between Feide and other federations.
- *LDAP:* Between the Feide login service and the host organizations.

Each interface is explained in the following sections. The interfaces are described according to the logical information flow. The physical message flow may differ somewhat; and this should be taken into account when considering security requirements for protecting the communication lines. The physical message flow is described in *Feide protocol description* [9].

5.1 Between end users and the feide login service

This primary function of this interface, which is based on web protocols, is to carry the end user's dialog with the Feide login service, called Moria. Moria provides information about the Feide service, and accepts the Feide name and password from the user for authentification purposes. This dialog is activated when the user does not yet have a running Feide session, or the service requires a fresh authentification (usually for security reasons).

The requirements to the dialog are specified in *Feide Login Service Requirement Specification* [1], chapter 5. Users shall be given the following information:

- that login is handled by the Feide login service,
- the name of the service that will be opened (it is not possible to log in to Feide in itself; login must be done through a service).
- which user attributes the service will receive through Feide,
- which measures are taken by Feide to protect information and privacy,
- that the user may select *not* to use single sign on (SSO); if SSO is deselected, every time a service requests authentication, a user dialog for supplying the user name and password will be activated,
- a link to the Feide web pages of frequently asked questions.

A technical specification of the information exchange between the user's web browser and the Feide login service is provided in *Feide protocol description* [9].

5.2 Between Feide services and the Feide login service

The interface is based on the web mechanisms for request redirection. The way these mechanisms are used are in agreement with SAML 2.0 [5]. The interface is defined in *Feide Login Service Requirement Specification* [1], chapter 3. The main responsibility of Feide is to confirm the authenticity of the user and to deliver those user attributes agreed upon in the contract between Feide and the service provider.

5.3 Between Feide and other federations

The detail specifications of this interface at protocol level may vary depending on the technology employed by the other federation. Neither end users, host organizations nor service providers are affected by the solutions chosen for this interface. Possible alternatives include:

- cross federation according to Liberty Alliance standards,
- web redirection modelled after the communication between service provider and Feide,
- LDAP lookup modelled after the communication between Feide and host organizations,
- other, possibly proprietary, authentication protocols, such as Shibboleth.

5.4 Between the Feide login service and host organizations

The interface is based on LDAP. The login service performs two operations against the host organization LDAP:

- The login service forwards the Feide name and password to the host organization, which confirms (or rejects) the password for this user
- The login service retrieves the user attributes for the now authenticated user.

The Feide architecture is designed to be independent of authentication method. The current release supports authentication by user name and password, transmitted across an encrypted connection only.

Details about the use of LDAP is found in *Feide protocol description* [9].

5.5 Between end users and Feide services.

This dialog does not involve Feide components. Data flows directly from one party to the other; Feide is not affected by, and has no access to, the information exchanged.

However, Feide authentication is essential for *establishing* the dialog. As the authentication mechanisms are based on web protocols, the dialog between the user and the Feide service will be based on web protocols as well. The Feide service may activate other protocols in the dialog with the user, but this is outside the scope of Feide.

A user connects to a *Feide service* rather than to Feide itself. Feide services not generally available will request login or authentication of the user before granting access. When a user establishes a connection, the Feide service may offer choices:

- local login independent of Feide,
- Feide login,
- authentication by other federations through Feide (cross federation),
- other external authentication services independent of Feide.

If the user is offered choices, the Feide login alternative should be identified by the Feide logo on the login page of the service. If the user selects this alternative, or cross federation through Feide, control is transferred to the Feide login service.

If authentication is *always* done by Feide, the service's login page is not required; control may immediately be transferred to the Feide login service before any interactive dialog is established between the user and the Feide service.

If the user is already logged in to Feide, that is, a Feide session is established, and both the user and the Feide service accepts single sign-on (SSO), the authentication is based on the established Feide session, and a second login is not required. A Feide service may, for security reasons, choose to always require a new login, even if the user has already established a Feide session.

When the service session is terminated, the Feide service may select to terminate the entire Feide session, or to terminate the service session only. Terminating the Feide session implies termination of *all* service sessions known to Feide for this user, including services other than the one requesting the logout. If one of the other services at some later time requests authentication, and there is no active Feide session, the user will be requested to log in anew.

Feide is not necessarily aware of all sessions between a user and a service, as Feide makes no requirement that all sessions are made known to the Feide login service. Feide logout only affects sessions associated with Feide authentication.

6 Host organizations

A host organization manages a group of Feide users: User name and password assignment, handling of attribute values for each user affiliated with the organization. Feide itself does not store neither authentication information nor user attributes, but assumes that the host organization will handle it.

The host organization is contractually obliged to fulfill a number of requirements to enable Feide to perform the authentication. Information about users must be stored in a secured system, and must be provided to Feide in a standard format. Otherwise, Feide could neither perform the authentication nor forward information about the users to service providers.

The LDAP directories of the host organizations may, as a whole, be viewed as a distributed database, with Feide as the central controller. This choice, rather than the alternative of a single, centralized database with all information managed by Feide, is supported by several arguments:

- Centralized data storage increases several risk factors, and will often increase the cost of storage facilities and of maintenance
- Updates in a centralized database cause more overhead. So, often update procedures are run less frequently, causing information to be less up to date, less reliable.
- Service providers should not have to relate directly to each individual host organization. They should be relieved of handling credentials such as passwords, but leave this to a centralized login facility.
- Correct use of user attributes depends on well defined semantics. This is ensured through the Feide LDAP scheme, defining value sets, syntax and semantics for all attributes, whether they are mandatory and whether attributes may be multivalued.

Because authentication directly depends on information from host organizations, keeping the information as up to date as possible is essential. In other words, all modifications, additions and deletions in the authoritative sources must be reflected in the data made available to Feide, and which may be forwarded by Feide.

This imposes the following requirements on a host organization:

- An identity management system, an IMS, must be established.
- The IMS handles information from authoritative sources. Handling routines for each source system may have to be modified to ensure that data of the highest possible quality is stored and delivered to the IMS. Quality assurance is a continuous activity.
- The roles and affiliations that user may have with the organization must be clearly defined. The organization's procedures must ensure consistent assignment and use of roles/affiliation values. Each Feide user must have one or more defined affiliations to his organization, selected from a small, but flexible, set of values. One affiliation value may be indicated as the primary one. Examples of affiliations are «student» or «employee».
- The basic data for the LDAP directory must be managed electronically. Common source systems are school management systems, employee registers and accounting systems.
- The organization must define one source system as *authoritative* for each person attribute: E.g. «name» and «phone number» has a single authority. All other occurrences of the attribute value are considered copies. If the correctness of the value is questioned, the value from the authoritative source takes presedence.

  If the organization has distinct groups of users, such as students and employees, the authoritative source may be different for each group, but for a given individual, the authoritative source must be unambiguously defined.
- The organization must offer an LDAP directory for the Feide login service.

- The LDAP directory must deliver information consistent with the Feide LDAP scheme.
- The LDAP directory tree must provide defined entry points available to the Feide login service.
- Information channels must be established to allow an automated and error free information flow between source systems and the LDAP directory interface to the Feide login service. This includes a mapping from the semantics of the source system data to the Feide LDAP scheme semantics. Automated procedures must make modifications of data in the source system available to Feide as soon as possible.

Becoming a host organization requires Feide approval. The following steps must be carried out:

- Install an IMS and an LDAP directory capable of delivering user attributes to Feide according to the stated specifications and requirements. Both open source systems and commercial systems from various vendors are available.
- Clean up the data in the source systems, e.g. remove outdated and duplicate entries, and verify that essential attributes are semantically consistent.
- Verify that procedures for management of personal data consistently ensures high data quality.
- Ensure that user password assignment procedures generate passwords of sufficient strength according to commonly accepted criteria. Ensure that all user names are unique and are assigned according to rules which will be unchanged for some time. An IMS may be of great help for these tasks.
- Verify that both authoritative data about each user, such as name and address, and generated data, such as user name and encrypted password, are available in the LDAP directory.
- Apply to Feide for becoming a host organization, and enclose documentation showing that the above points are taken care of.
- Sign a contract with Feide.

  If the organization also plans to offer Feide services, this is covered by the same contract.
- Make the LDAP directory available to the Feides login service.

7 Users

Feide makes requirements to users indirectly only: Formal agreements are made between Feide and the user's host organization, and the user answers to his host organization.

The Feide architecture and design is modelled to protect personal information from misuse and for granting access for users to their own data, the right to be informed about which attributes Feide services will have access, and the right to sanction attribute forwarding to services.

The host organization defines and enforces rules for use of the IT systems. Feide requires the organization to maintain well defined rules for management of personal information.

The guidelines should at least cover:

- who will be assigned (electronic) identities in the IT infrastructure of the organization (e.g. students, teachers, other employees, others affiliated with the organization),
- how the identity is established,
- how the identity is maintained,
- permitted use of the identity,
- when the identity is to be terminated.

Feide names may, and will, be used outside the host organization. Feide requires the host organization to hold their users responsible for adhering to guidelines for appropriate use of IT equipment. The guidelines are defined by the organization; Feide may provide assistance when establishing or updating them.

Feide offers a standard template for such guidelines in the educational sector. The template consists of a general part, common elements determined by the type of organization, as well as specific elements for each individual organization.

The host organization is responsible for, at a sufficient level of confidence, correct identification of persons to be considered users affiliated with the organization. If a person provides incorrect information to the organization, e.g. by using forged identity documents, this is an issue between the host organization and the person; the organization is responsible for handling or preventing this situation. Feide assumes unconditionally that the information in the LDAP directory is correct and valid, in accordance with the agreements and contracts applying to Feide's use of the LDAP directory.

8 Feide services

A service is an offer to perform a defined task for a user, such as perform a database search, create a printout, transfer and email or display information and learning material. Feide services may be *local*, available only to users in the same organization as the service provider (e.g. a university offering services to their own students), or *common services* which may be made available even to Feide authenticated users in other organizations.

For common services, the service provider may further restrict availability e.g. to users who has an established customer relationship. This is particularly relevant for commercial services. Any access restrictions other that that of organizational affiliation must be enforced by the service itself.

8.1 Establishing a Feide service

A service authenticating users through Feide is a Feide service. To become a Feide service, the service provider, i.e. the «owner» of the service, must be approved by Feide. The approval process is initiated through an application to Feide, and leads up to a contract for a service that communicates with the Feide login service for authentication and user attribute access. The contract sets the following conditions:

- which user attributes will be delivered to the service,
- contracts for technical and administrative issues, as well as a list of persons entitled to negotiate changes in the contract between the service provider and Feide, and to negotiate changes in the configuration data used by Feide for the service,
- procedures for reporting modifications to the service or the addition of new common services,
- whether the service will use single sign-on or not,
- whether the service will use one-time federation, federation at first use or provisioning - in other words, whether the service will be of category A, B or C (see below),
- requirements for handling of personal information.

8.2 Protocols

Feide services are offered to users through a web interface. The communication runs directly between service and user; only in the authentication phase, the user is redirected to the Feide login service.

During authentication, services communicate with Feide through protocols defined by Liberty Alliance: SAML 2.0 and ID-WSF [5]. The requirements for this interface is defined in *Feide Login Service Requirement Specification* [1], chapter 3. The service provider may use any software that fulfills these requirements.

8.3 Service categories

Feide services are classified into three categories by how the service relates to users:

A) The service is not aware of the identity of the user. The user is anonymous or known only by general properties such as «student», «employee» etc.

B) A service account is established the first time the user contacts the service. The service may make user specific adaptations, such as managing a «user profile» that may be stored permanently.

C) Information about a group of users may be uploaded to a service provider, and all these users will be predefined in the service provider's user directory. The uploading process is termed «provisioning». Typical groups of new users are e.g. all new students entering the university at the start of a semester.

8.4 Sample implementation for service providers

Feide offers a standard template for a SAML 2.0 service provider which may simplify adaptation of services to Feide.

The sample application template may be used as a basis when judged as useful by the developer of the service; there are no restrictions on use of tailor made, proprietary software or other third party libraries. Solutions based on tailor made software may allow more control of the authentication process, but requires detail knowledge of the protocol standards and the profile selected by Feide.

9 Feide login

Feide Login Service Requirement Specification [1] describes all direct and indirect requirements to this Feide component, several of which have been discussed earlier in this document.

9.1 Survey of the login sequence

- When a service makes a request to a Feide service at which the user already has an active and valid service session, the Feide login service is not consulted; the request is honored with no further authentication.
- If the user does not have an active service session, the Feide service temporarily holds back the response, and transfers control to the Feide login service for authentification.
- If the user is already logged in to Feide, the Feide service receives a confirmation of the authentication, together with relevant attribute values for the user, and no user dialog with Feide is opened. The Feide service then processes the original user request.
- If the user is not logged in to Feide, but requests authentication through another federation (i.e. cross federation), Feide transfers control to the other federation, which (unless the user has already been authenticated) will open an interactive user dialog. Feide receives an authentication confirmation, and forwards this to the Feide service, which may then process the original user request.
- If the user is not logged in to Feide, and has selected Feide authentication, the Feide login service opens an interactive user dialog for obtaining the user name and password. The login service forwards the user name and password to the LDAP directory of the user's host organization. Upon successful authentication, the login service returns a confirmation to the Feide service, along with a selection of user attributes according to the service contract. The service may then process the original user request.

Remarks:

- Management of service sessions, that is, handling information about which users are connected to the Feide service at any given time, is performed outside the scope of Feide. The standard software delivery from Feide offers tools for these tasks, which may or may not be used by the service provider.
- The Feide service may request Feide to open an authentication dialog with the user, even if the user already has an active and valid Feide session. This is used by services requiring a higher security level, to prevent a workstation logged into Feide, but temporarily left by the user, from being hijacked by intruders.

Appendix A: Referanser

[1] Feide Login Service Requirement Specification, version 2005-11: http://www.feide.no/dokumenter/feide-login-service-requirementspecs-ver200511.pdf

[2] norEdu LDAP Schema, version 1.4: http://feide.no/dokumenter/norEdu-1.4Final.pdf

[3] eduPerson/eduOrg LDAP Schema, http://www.educause.edu/eduperson

[4] SAML Specifications, OASIS Security Services Technical Community, http://www.oasis-open.org/committees/security/

[5] Liberty Alliance Project, Identity Federation Framework, http://www.projectliberty.org

[6] Lov om behandling av personopplysninger (The Norwegian Privacy Protection Act), http://www.lovdata.no/all/nl-20000414-031.html

[7] Shibboleth, http://shibboleth.internet2.edu/

[8] Feide risk analysis [internal report].

[9] Feide protocol description [not yet available]

[10] Moria Web Authentication Service, http://moria.sourceforge.net